

A DEFLATED SCHUR COMPLEMENT METHOD FOR THE ITERATIVE SOLUTION OF A HIGH-ORDER DISCONTINUOUS ELEMENT DISCRETIZATION OF THE POISSON EQUATION

SUMEDH JOSHI

*Center for Applied Mathematics
657 Rhodes Hall, Cornell University
Ithaca NY, 14850*

PETER DIAMESSIS

*School of Civil and Environmental Engineering
105 Hollister Hall, Cornell University
Ithaca NY, 14850*

ABSTRACT. A combination of block-Jacobi and deflation preconditioning is used to solve a high-order discontinuous element-based collocation discretization of the Schur complement of the Poisson-Neumann system as arises in the operator splitting of the incompressible Navier-Stokes equations. The ill-posedness of the Poisson-Neumann system manifests as an inconsistency of the Schur complement problem, but it is shown that this can be accounted for with appropriate projections out of the null space of the Schur complement matrix without affecting the accuracy of the solution. The block-Jacobi preconditioner, combined with deflation, is shown to yield GMRES convergence independent of the polynomial order of expansion within an element. Finally, while the number of GMRES iterations does grow as the element size is reduced (e.g. h -refinement), the dependence is very mild; the number of GMRES iterations roughly doubles as the element size is divided by a factor of six. In light of these numerical results, the deflated Schur complement approach seems practicable, especially for high-order methods given its convergence independent of polynomial order.

1. INTRODUCTION

1.1. Background. Domain decomposition methods have been recently applied to high-order, discontinuous, discretizations of elliptic problems with good success [3, 2, 4, 5, 19]. These methods aim to build preconditioning strategies for iterative Krylov-type solvers to obtain an accurate solution whose convergence is independent of the parameters of the grid. These approaches have largely focused on the widely-used discontinuous Galerkin class of numerical methods, and many leverage the symmetric positive-definite nature of the DG discretization.

Analogous to the DG class of discretization methods, the Spectral Multidomain Penalty Method (SMPM) is a high-order discontinuous variant of the spectral element method that uses spectral differentiation matrices to compute derivatives [16]. Because spectral differentiation matrices are themselves unsymmetric [7] the operator matrices resulting from SMPM are unsymmetric and not self-adjoint. Nevertheless, the SMPM discretization has been used to solve complex, large-scale, environmental fluid mechanics problems on hundreds of processors and with hundreds of millions of unknowns [10, 1, 11].




FIGURE 1. A depiction of the logical arrangement of a 2×2 element spectral multi-domain penalty method (SMPM) grid with 10×10 Gauss-Lobatto-Legendre points in each element denoted V_j for $j = 1, 2, 3, 4$. The inter-element continuity fluxes are represented with R_{ij} with $i \neq j$.

1.2. Spectral multi-domain penalty method. In this work, we describe a method for solving the 2D Poisson-Neumann system that arises within the time-splitting of the 2D incompressible Navier-Stokes equations [13], and is given on a domain $\Omega \subset \mathbb{R}^2$ as

$$(1) \quad \begin{aligned} \nabla^2 u &= f \text{ on } \Omega \\ n \cdot \nabla u &= g \text{ on } \partial\Omega. \end{aligned}$$

Ω is discretized into an $m_y \times m_x$ cartesian quadrilateral element grid with elements Ω_{ij} , where $i = \{1, \dots, m_y\}$ and $j = \{1, \dots, m_x\}$. Within each element is a 2D Gauss-Lobatto-Legendre (GLL) grid with n GLL points per direction for a total of n^2 grid points per element. Since the discretization is discontinuous, function values are allowed to differ along the $2n$ grid points on the boundary of each pair of elements, and thus the full grid has a total of $r = n^2 m_x m_y$ grid points.

Now we define the SMPM element matrices and inter-element continuity conditions. Let $Lu = f$ represent the discrete Poisson-Neumann system on $\Omega \subset \mathbb{R}^2$ a domain discretized into an $m_x \times m_y$ element mesh with each element V_i smoothly and invertibly mapped from the master element $[-1, 1] \times [-1, 1]$. On each element a two-dimensional Gauss-Lobatto-Legendre (GLL) grid with n points in each direction is constructed and used to evaluate the Lagrange interpolant basis and their derivatives by way of spectral differentiation matrices[7]. Thus each element contains n^2 grid points. If V_i and V_j share the n GLL points along one of their four boundaries, then each element owns a copy of those n GLL nodes in order maintain the discontinuous nature of this method. Thus as a matrix, $L \in \mathbb{R}^{r \times r}$ is of dimension $r = n^2 m_x m_y$, where r denotes the total number of nodes in the grid Ω .

In the SMPM the weak inter-element continuity condition is of Robin type, and is enforced by the flux $R_{ij} : \partial V_j \rightarrow \partial V_i$ from element V_j into V_i for V_i, V_j with an adjacent boundary $\partial V_j \cap \partial V_i$ consisting of n grid points. R_{ij} is defined as

$$(2) \quad R_{ij} = I + \hat{n}_i \cdot \nabla$$

where $\hat{n}_i : \partial V_i \rightarrow \mathbb{R}^2$ is the outward pointing normal vector of ∂V_i and I is the identity operator. A depiction of a 2×2 element grid with the inter-element fluxes is shown in Fig. 1, in which the elements V_1, V_2, V_3, V_4 have been separated to emphasize the discontinuous nature of the SMPM.

The physical boundary conditions are Neumann, and are given on $\partial V_i \cap \partial\Omega$ as $n_i \cdot \nabla$ where n_i is again the outward pointing normal vector. Given a function u , on an element V_i the residual in the spectral multi domain penalty method is given by the sum of the Laplacian, the inter-element continuity mismatch, and

the boundary condition mismatch as

$$(3) \quad L_i u_i = \nabla^2 u_i + \tau_i \left(R_{ii} u_i - \sum_{j \in N(i)} R_{ij} u_j \Big|_{\partial V_i \cap \partial V_j} \right) + \tau_i \hat{n}_i \cdot \nabla u_i \Big|_{\partial V_i \cap \partial \Omega} = f_i + \tau_i g_i.$$

Here, g_i is the boundary value of the Neumann boundary condition restricted to element V_i , and $N(i)$ is the index set of elements adjacent to V_i . The inter-element continuity, external boundary conditions, and the PDE are all satisfied weakly, since the residual is the sum of these three components. The penalty parameter τ_i represents the degree to which the inter-element continuity and boundary conditions are weighted in the residual relative to the PDE, and the optimal choice of τ_i is determined by stability criteria for hyperbolic problems[15, 16], and a heuristic for the Poisson problem [13].

1.3. Construction of the Schur complement problem. As shown in Figure 2 the domain Ω is discretized with a collection of *elements* V_j , each invertibly mapped from the master element $[-1, 1] \times [-1, 1]$. Ω is decomposed into $m_x m_y$ many *sub-domains* Ω_i , with each sub-domain corresponding to a single element in the mesh. Since the SMPM is a high-order method, each element represents large, local, and dense linear algebraic operations. Thus, it makes sense from a domain decomposition perspective for the elements and sub-domains to coincide.

Along each of the $(m_x - 1)(m_y - 1)$ interfaces between the sub-domains/elements are $2n$ GLL nodes (n nodes on either side of each interface). Denote as k the number of interfacial nodes in the domain decomposition, and this set of k interface nodes as Γ . The discrete Poisson operator L (Eq. 3) is decomposed into a local term and an inter-subdomain flux term which is used to construct the Schur problem. This operator decomposition comprises three operators which are defined below.

First, denote as $E : \Gamma \rightarrow \Omega$ the inclusion map that maps from the interfacial grid Γ to Ω . As a matrix, $E \in \mathbb{R}^{r \times k}$ and is composed of zeros and ones, E^T is the restriction from the full grid to the interface grid Γ , and $E^T E = I \in \mathbb{R}^{k \times k}$ the identity matrix. Naturally EE^T is not an identity matrix.

Second, define an operator $B : \Omega \rightarrow \Gamma$ that consists of the inter-subdomain Robin boundary fluxes. B represents all of the the inter-element fluxes R_{ij} . As a matrix, $B \in \mathbb{R}^{k \times r}$, since it computes $I + \hat{n} \cdot \nabla$ within a subdomain using spectral differentiation matrices and assigns it to the interface of its neighbor.

Finally, define the operator $A : \Omega \rightarrow \Omega$, which represents the part of L that is entirely local to one subdomain. A consists of the Laplacian part of L , the boundary condition mismatch, and the R_{ii} terms in Eq. (3). Since A is entirely local to each subdomain, as a matrix $A \in \mathbb{R}^{r \times r}$ is block-diagonal. A represents $m_x m_y$ decoupled homogenous Poisson-Robin boundary value problems, and as such is invertible and block diagonal.

These three operators are defined so that their combination yields the SMPM Poisson-Neumann operator

$$(4) \quad Lu = Au + EBu = f.$$

Notice that the action of EB couples the subdomains only weakly since its action is combined with A in the residual $(A + EB)u = f$. This weak enforcement of inter-subdomain continuity in the SMPM allows for decoupling the subdomains by decoupling the action of B from that of A . To accomplish this, we seek a vector $v \in \mathbb{R}^k$ on the interfacial nodes, $\bigcap_i \partial \Omega_i$, such that the solution to $A_i u_i = f_i - E v_i$ on each Ω_i subdomain also solves $Lu = f$. By writing

$$(5) \quad Au = f - EBu$$

it is clear that $v = Bu$, the image of the solution under the inter-subdomain flux operator. Because B is a contraction (mapping from the full grid to the interfacial grid), finding the image Bu is easier than finding u itself; its value is given by the solution to the system

$$(6) \quad \begin{bmatrix} A & E \\ B & -I \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix}.$$

As is evident, any $[u, v]^T$ that solves this system also solves $Lu = f$, and satisfies $v = Bu$; this system represents splitting the range of A and B in obtaining a solution of L . Taking one step of block Gaussian

elimination of A in this matrix to result in the upper triangular system

$$(7) \quad \begin{bmatrix} I & A^{-1}E \\ 0 & -I - BA^{-1}E \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} A^{-1}f \\ -BA^{-1}f \end{bmatrix},$$

we then obtain v as the solution to the following system,

$$(8) \quad (I + BA^{-1}E)v = BA^{-1}f,$$

which represents the Schur complement system of A in Eq. (6). A back-substitution of v into Eq. (7),

$$(9) \quad u = A^{-1}(f - Ev),$$

results in u the solution of $Lu = f$. Because the SMPM is a discontinuous element discretization, A is block diagonal and invertible, and so all divisions of A are easily parallelized; the expensive part of the above is obtaining the solution of the Schur complement system. In the rest of this paper, the focus is on efficiently obtaining this solution.

1.4. Inconsistency of the Poisson-Neumann system. Prior to obtaining the solution to Schur complement system, there remains the important point of dealing with the rank-deficiency of the Poisson-Neumann operator L . The Poisson-Neumann equation is ill-posed in the continuous sense, and so the SMPM operator L is rank-deficient and has non-trivial left and right null spaces of dimension one. In symmetric discretizations, the kernel vector is the constant vector, but since L is unsymmetric its left and right null spaces are different and only the right null space is constant vector. To ensure consistency and solvability the right-hand-side vector f is projected out of the left null space of L [22] and instead of $Lu = f$, the regularized system solved is

$$(10) \quad Lu = \tilde{f}$$

where $\tilde{f} = f - u_L u_L^T f$ is f projected onto the range space of L and $u_L \in \mathbb{R}^r$ is the unique vector with unit norm that satisfies $\|u_L^T L\|_2 = 0$. The solution u then is only known up to an indeterminant additive constant vector. The rank deficiency of L is inherited by the Schur complement system, and thus another regularization is required to project the Schur right hand side $b_S = BA^{-1}\tilde{f}$ out of the left null space of the Schur complement system. Thus the Schur complement system,

$$(11) \quad Sx_S = b_S,$$

is modified to read

$$(12) \quad Sx_S = b_S - u_S u_S^T b_S.$$

To summarize, the method for obtaining the solution u to $Lu = f$ is shown in Algorithm 1. The statement $\text{GMRES}(S, b_S)$ in Step 4 is meant to represent the solution of the linear system $Sx_S = b_S$ with the Generalized Minimum Residual Method (GMRES).

Algorithm 1 Schur complement method with null space projections.

Input: f, u_L, u_S

Output: u

- 1: $f \leftarrow f - u_L u_L^T f$
 - 2: $b_S := BA^{-1}f$
 - 3: $b_S \leftarrow b_S - u_S (u_S^T b_S)$
 - 4: $x_S := \text{GMRES}(S, b_S)$
 - 5: $u \leftarrow A^{-1}(f - Ex_S)$
-




FIGURE 2. A sample domain with $m_x = 4$ and $m_y = 4$ elements in x and y respectively. The elements have been separated to emphasize the discontinuous nature of the discretization. The interfaces that together form the Schur grid Γ are highlighted in color, with the colors signifying the blocks of the block-Jacobi preconditioner. Interior blocks of the block-Jacobi preconditioner consist of eight interfaces; exterior blocks consist of six or four interfaces.

2. DEFLATED SCHUR COMPLEMENT METHOD

The iterative solution of the Schur complement system with GMRES requires an efficient preconditioner. Many preconditioning techniques for the Schur complement system have been proposed [23, 9, 6, 24] with most relying on two-level preconditioners: a local preconditioner that can be applied in parallel and a coarse global preconditioner to speed across-grid communication of components of the residual. An example is the two-level additive-Schwarz preconditioner in which overlapping block-diagonal components are solved in parallel, augmented with a coarse grid correction to communicate information across the grid [14, 13]). In this work, a non-overlapping block-diagonal/block-Jacobi preconditioner is used, augmented with deflation, to achieve Krylov subspace convergence rates independent of the polynomial order, p , and weakly dependent on the element size, h .

First, note that for sparse matrices with non-zeros clustered around the diagonal, computing the inverse of blocks along the diagonal separately can be a useful preconditioning technique. Block-Jacobi preconditioners have been shown to be effective for the Schur complement of elliptic operators [8], especially when combined with coarsened-grid preconditioners [18, 21, 20]. Here, for preconditioning the Schur matrix, S , a block-Jacobi preconditioner is assembled in which a single block represents the coupling between the four interfaces bounding one element, and their corresponding interfaces in neighboring elements (for a total of $8n$ grid points in one block). The elements in the $m_y \times m_x$ grid are divided in a checkerboard pattern as can be seen in Fig. 2, in which adjacent nodes are grouped together in blocks by color. The colors correspond to the blocks of the block diagonal preconditioner for the Schur complement matrix. Denoting the block-Jacobi preconditioner matrix as M , the preconditioned Schur complement system that is solved with GMRES is

$$(13) \quad SM^{-1}x'_S = b_S$$

and the solution is obtained by a final division by M

$$(14) \quad x_S = M^{-1}x'_S.$$

Since M is explicitly block-diagonal (i.e. any non-zeros of S coupling the blocks of M are ignored in the factorization of M), divisions by M can be computed efficiently in parallel. The algorithmic summary of the preconditioned Schur complement method is given in Algorithm 2.

Algorithm 2 Preconditioned Schur complement method

Input: b, u_L, u_S, u_C

Output: x

- 1: $b \leftarrow b - u_L u_L^T b$
 - 2: $b_S := BA^{-1}b$
 - 3: $b_S \leftarrow b_S - u_S(u_S^T b_S)$
 - 4: $x := \text{GMRES}(SM^{-1}, b_S)$
 - 5: $x \leftarrow M^{-1}x$
 - 6: $x \leftarrow A^{-1}(b - Ex)$
-

Working in tandem with other preconditioners, deflation methods aim to accelerate the convergence of Krylov methods by eliminating (or “deflating”) components of the residual within a chosen subspace. The subspace is usually chosen to be a span of approximate eigenvectors of the operator corresponding to slowly converging eigenvalues. Thus, the problematic eigenvalues are solved directly using a coarsened version of the operator, and the remaining components of the residual are eliminated by a Krylov solver. Here, to augment the block-Jacobi preconditioner described in the previous section, a deflation method is used as the coarse-grid correction method, following the procedure in Ref. [12]. The deflation vectors are chosen to be a set of d column vectors $Z \in \mathbb{R}^{k \times d}$ where $d \ll k$, and $k = \dim(S)$. These deflation vectors are chosen to be discrete indicator vectors, equal to 1 on each pair of interfaces between two elements and zero everywhere else. Denoting as Γ_j an interface between two elements, the i -th entry in the j -th deflation vector is given by

$$(15) \quad (z_j)_i = \begin{cases} 1 & : \text{if } x_i \in \Gamma_j \\ 0 & : \text{if } x_i \notin \Gamma_j \end{cases},$$

thus each vector is active on one pair of interfaces in the Schur grid. The matrix of these vectors $Z = [z_1, z_2, \dots, z_d]$ defines a coarse version of the Schur problem, $C = Z^T S Z \in \mathbb{R}^{d \times d}$, and two projections

$$(16) \quad P = I - SZC^{-1}Z^T$$

$$(17) \quad Q = I - ZC^{-1}Z^T S$$

each of size $\mathbb{R}^{k \times k}$. As a matrix $Z^T \in \mathbb{R}^{d \times k}$ is a contraction operator that maps grid functions on the Schur grid to the coarse grid, and its transpose is a prolongation operator. The intuition behind the projections P and Q is that they project out of the subspace on which $ZC^{-1}Z^T$ is a good approximation of the left (in the case of Q) or right (in the case of P) inverse of S . Thus the projections map onto the complement of the subspace on which the coarse matrix C approximates the Schur matrix S well. Finally, note that all applications of C^{-1} require their own regularization since C inherits rank deficiency from S ; denote as u_C the left null space of C in the following. Deflation proceeds by noting that the solution of the Schur complement problem $Sx_S = b_S$ can be decomposed into

$$(18) \quad x_S = (I - Q)x_S + Qx_S.$$

Then, the first term is just $ZC \setminus (Z^T - u_C u_C^T Z^T)(b_S - u_S u_S^T b_S)$, which can be computed directly since C is small. The second term is obtained by way of GMRES on the deflated and right-preconditioned system $PSM^{-1}x_S = P(b_S - u_S u_S^T b_S)$ and then post-multiplying by Q , finally assembling the solution as

$$(19) \quad \begin{aligned} x_S = & ZC \setminus (Z^T - u_C u_C^T Z^T)(b_S - u_S u_S^T b_S) \\ & + QM^{-1}\text{GMRES}(PSM^{-1}, P(b_S - u_S u_S^T b_S)). \end{aligned}$$

Because P projects out of the coarse space, the GMRES solution of $PSM^{-1}x_S = P(b_S - u_S u_S^T b_S)$ minimizes only the component of the residual that cannot be well-approximated by the coarse solution. This formulation of deflation-augmented right-preconditioning is an extension of the work in Ref. [12] to a rank-deficient matrix. For completeness, Algorithm 3 depicts the algorithmic summary of the deflation method in which the notation $\text{GMRES}(A, b)$ is intended to represent the solution of a linear system $Ax = b$ with GMRES.

Algorithm 3 Deflated and preconditioned Schur complement method

Input: b, u_L, u_S, u_C
Output: x

- 1: $\tilde{b} \leftarrow b - u_L u_L^T b$
 - 2: $\tilde{b}_S := B A^{-1} \tilde{b}$
 - 3: $\tilde{b}_S \leftarrow \tilde{b}_S - u_S (u_S^T \tilde{b}_S)$
 - 4: $x_1 := \text{GMRES}(P S M^{-1}, P \tilde{b}_S)$
 - 5: $x_1 \leftarrow Q M^{-1} x_1$
 - 6: $x_2 := Z^T \tilde{b}_S - u_C u_C^T Z^T \tilde{b}_S$
 - 7: $x_2 \leftarrow Z C \setminus x_2$
 - 8: $x := x_1 + x_2$
 - 9: $x \leftarrow A^{-1}(b - E x)$
-

3. PERFORMANCE




FIGURE 3. *Left:* The error decay as the initial grid with $p = m_x = m_y = 4$ is refined in both h and p . Notice the decay of error is polynomial in h but exponential in p . *Middle:* h refinement of the initial grid in which $p = m_x = m_y = 4$, with the analytic L^2 error plotted against the number of GMRES iterations required to achieve a relative tolerance of 10^{-10} . h is gradually reduced as m_x and m_y are increased from 4 to 32. *Right:* p refinement of the initial grid in which $p = m_x = m_y = 4$, with the analytic L^2 error plotted against the number of GMRES iterations required to achieve a relative tolerance of 10^{-10} . p is gradually increased from 4 to 24.

To study the performance of the deflated Schur complement method, the domain $\Omega = [0, L_x] \times [0, L_y]$ is discretized with $n = 5$ (4th order polynomials), $m_x = 4$ and $m_y = 4$. On this grid the following Poisson problem was solved,

$$(20) \quad \begin{aligned} \nabla^2 u &= \cos(\lambda \pi x / L_x) \cos(\lambda \pi y / L_y) \\ n \cdot \nabla u &= 0, \end{aligned}$$

with $\lambda \in \mathbb{N}$, which has the analytic solution

$$(21) \quad u_a(x, y) = -\frac{L_x, L_y}{2\lambda^2 \pi^2} \cos(\lambda \pi x / L_x) \cos(\lambda \pi y / L_y).$$

Besides evaluating the performance of the deflated Schur complement method with respect to GMRES convergence, comparing against an analytic solution makes it clear whether the discretization is exhibiting spectral convergence with respect to a known solution. In the following subsections, we evaluate the GMRES convergence properties of the deflated Schur complement method as the initial grid is refined in both p , the

polynomial order, and h , the element edge size (L_x/m_x). The error e is computed as the L^2 norm difference of the numerical solution, u , against the analytic solution, u_a ,

$$(22) \quad e = \|u_a - u\|_2.$$

In all of the examples below, $\lambda = 7$.

3.1. Spectral convergence. First, note that as shown in Fig. 3(a), the Schur complement approach to solving the SMPM discretization converges to the true solution as a polynomial of h and as an exponential of p . This is made clear by the relatively gradual decay of the error in h as compared with that in p , and is a hallmark of high-order methods like the spectral multi domain penalty method [13]. For the smallest values of h and largest values of p the error does not decay monotonically to zero; this may be an artifact of the ill-conditioning of the spectral differentiation matrices, which only worsens as p grows and h decreases.

3.2. h -refinement. Having established that the Schur complement method converges to the analytic solution at a rate equal to a polynomial of h , we now examine how the convergence of GMRES is affected as h is refined. Starting again with a grid $p = m_x = m_y = 4$ on a domain $\Omega = [0, L_x] \times [0, L_y]$, the number of elements m_x and m_y is iteratively grown, yielding a refinement in $h = L_x/m_x$. The results of this refinement study are shown in Fig. 3(b) for values of $m_x, m_y = \{4, 8, 12, 16, 20, 24, 28, 32\}$. The Schur complement is assembled exactly as before, with decomposition along all internal element boundaries, and so the dimension of the Schur complement problem grows as $\mathcal{O}(m_x m_y)$.

The GMRES algorithm is employed to reduce the residual to a relative tolerance of 10^{-10} , at which point the rest of the Schur complement algorithm is employed to reconstruct the full solution (c.f. Step 5 in Algorithm 1). The number of GMRES iterations required to achieve this tolerance is depicted along the horizontal axis in Fig. 3(b). In the vertical axis of the same figure is shown the analytic error $\|u_h - u_a\|_2$ in the resulting solution. In all cases, the number of GMRES iterations grows as h is refined. This is first because the conditioning of the element stiffness matrices degrades as the element size goes to zero [7], and second because the dimension of the Schur complement matrix grows as m_x and m_y grow.

However, deflating the Schur complement in the GMRES solver tempers the growth of the number of iterations as h is refined. At the finest grid, in which $m_x = m_y = 32$, the unpreconditioned Schur complement method takes nearly 110 Krylov iterations to converge; the deflated Schur complement method takes just over 30. Furthermore, asymptotically, the number of iterations required to obtain a solution grows much more mildly in the deflated Schur complement approach. Thus, it is observed that deflation, while not eliminating the dependence on h , strongly mitigates the growth in GMRES iterations as h is reduced. Finally, note that all of the Schur complement approaches shown in Fig. 3(b) are very efficient. For example, the number of grid points r in the systems solved in Fig. 3(b) is $r = n^2 m_x m_y$. Since $n = 5$ and m_x and m_y grow to be 32 each, the total number of grid points grows to over 25,000, which is greater by several orders of magnitude than the number of GMRES iterations required to obtain a solution that is correct to ten decimal places.

3.3. p -refinement. The results of a refinement study in p are shown in Fig. 3(c) for values of $p = \{4, 6, 8, 10, 12, 14\}$, and $m_x = m_y = 4$. As p is increased, the size of the Schur complement matrix grows as p^2 , and its conditioning properties worsen due to the h^{-p} conditioning of the spectral differentiation matrices embedded within it [7]. Nevertheless, it is observed in Fig. 3(c) that the convergence properties of GMRES are essentially unaffected by p -refinement when the block-diagonal preconditioner is applied to the Schur complement matrix. Even the unpreconditioned Schur complement method only shows mild growth in GMRES iteration count as p grows. While it may seem that this p -independent convergence may depend on smoothness of the right-hand-side, this result has been confirmed for random white-noise right-hand-side vectors as well [17]. It appears that GMRES convergence of the preconditioned Schur complement method is robust to refinements in p , which is particularly useful given that the error in the solution decays exponentially with p . Finally, note again that all variants of the Schur complement method shown in Fig. 3(c) are exceedingly efficient. Since the grid grows as p^2 , for the largest value of p , the number of grid points is $r = 3600$; yet, GMRES converges to ten digits of accuracy in less than 40 iterations in all cases, and in under 15 iterations in the preconditioned/deflated cases.

4. CONCLUSION

A preconditioned Schur complement technique for solving the spectral multi-domain penalty method discretization of the Poisson-Neumann system was developed. The preconditioning method relies on a local block-Jacobi preconditioner and subspace deflation to solve a coarse component of the residual. By using both a local (block-Jacobi) and global (deflation) preconditioner, convergence of GMRES only mildly dependent on the grid resolution, h , and independent of polynomial order, p , is possible. Since the error in the SMPM decays exponentially with p , achieving GMRES convergence independent of p is very useful in practice, as it allows for high-accuracy solutions by increasing p at minimal additional cost.

These ideas have already been leveraged within a high-order incompressible Navier-Stokes solver; the Poisson-Neumann system arises there as part of the incompressibility constraint on the flow. By leveraging the deflated Schur complement algorithm described herein, problems on grids with millions of unknowns are solved with $\mathcal{O}(10)$ GMRES iterations. The scalable efficient implementation of this deflated Schur complement approach is currently underway, but these methods show promise for parallelization since they are inspired from domain decomposition methods which are designed, in a sense, for distributed-memory parallel computing.

A natural extension of this work is to generalize to three-dimensional problems. We have shown already that if a periodic third dimension can be assumed, a Schur factorization can be used to efficiently extend these preconditioned Schur methods to three-dimensional problems [17]. However, for more general three-dimensional approaches (in which periodicity cannot be assumed) more work is required to demonstrate the efficacy of the Schur complement method espoused in this work. That having been said, there are few theoretical hurdles for such an extension. In particular, some Schur complement methods are plagued with so-called *cross-points*, or grid points belonging to multiple elements, whose treatment requires special care in constructing the Schur complement problem. Because the SMPM is a discontinuous discretization technique, there are no cross-points; every grid point belongs to only one element. Thus, the extension of the Schur complement method to three dimensions is far less cumbersome in the SMPM than in continuous element methods.

REFERENCES

- [1] Ammar M. Abdilghanie and Peter J. Diamessis. The internal gravity wave field emitted by a stably stratified turbulent wake. *Journal of Fluid Mechanics*, 720:104–139, 2013.
- [2] Paola F. Antonietti and Blanca Ayuso. Schwarz domain decomposition preconditioners for discontinuous Galerkin approximations of elliptic problems: non-overlapping case. *ESAIM: Mathematical modeling and numerical analysis*, 41(1), 2007.
- [3] Paola F. Antonietti and Paul Houston. A Class of Domain Decomposition Preconditioners for hp-Discontinuous Galerkin Finite Element Methods. *Journal of Scientific Computing*, 46:124–149, 2010.
- [4] Kolja Brix, Martin Campos Pinto, Claudio Canuto, and Wolfgang Dahmen. Multilevel preconditioning of discontinuous Galerkin spectral element methods. Part I: geometrically conforming meshes. *IMA Journal of Numerical Analysis*, 35:1487–1532, 2015.
- [5] Claudio Canuto, L. F. Pavarino, and a. B. Pieri. BDDC preconditioners for continuous and discontinuous Galerkin methods using spectral/hp elements with variable local polynomial degree. *IMA Journal of Numerical Analysis*, 34:879–903, 2013.
- [6] L. M. Carvalho, L. Giraud, and P. Le Tallec. Algebraic Two-Level Preconditioners for the Schur Complement Method. *SIAM Journal on Scientific Computing*, 22:1987–2005, 2001.
- [7] Bruno Costa and Wai Sun Don. On the computation of high order pseudospectral derivatives. *Applied Numerical Mathematics*, 33:151–159, 2000.
- [8] W Couzy and M O Deville. A Fast Schur Complement Method for the Spectral Element Discretization of the Incompressible Navier-Stokes Equations. *J. Comput. Phys.*, 116:135–142, 1995.
- [9] Jean-Michel Cros. A preconditioner for the Schur complement domain decomposition method. *14th International Conference on Domain Decomposition Methods*, pages 373–380, 2002.
- [10] Peter J. Diamessis and Larry G. Redekopp. Numerical Investigation of Solitary Internal Wave-Induced Global Instability in Shallow Water Benthic Boundary Layers. *Journal of Physical Oceanography*, 36:784–812, 2006.

- [11] Peter J. Diamessis, Geoffrey R. Spedding, and J. Andrzej Domaradzki. Similarity scaling and vorticity structure in high-Reynolds-number stably stratified turbulent wakes. *Journal of Fluid Mechanics*, 671:52–95, 2011.
- [12] Yogi a. Erlangga and Reinhard Nabben. Deflation and Balancing Preconditioners for Krylov Subspace Methods Applied to Nonsymmetric Matrices. *SIAM Journal on Matrix Analysis and Applications*, 30:684–699, 2008.
- [13] J.A. Escobar-Vargas, P.J. Diamessis, and T. Sakai. A spectral quadrilateral multidomain penalty method model for high Reynolds number incompressible stratified flows. *International Journal for Numerical Methods in Fluids*, 75(March):403–425, 2014.
- [14] Paul Fischer. An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations. *Journal of Computational Physics*, pages 1–35, 1997.
- [15] J. S. Hesthaven. A Stable Penalty Method for the Compressible Navier–Stokes Equations: II. One-Dimensional Domain Decomposition Schemes. *SIAM Journal on Scientific Computing*, 18(3):658–685, 1997.
- [16] J. S. Hesthaven. A Stable Penalty Method for the Compressible Navier–Stokes Equations: III. Multidimensional Domain Decomposition Schemes. *SIAM Journal on Scientific Computing*, 20(1):62–93, 1998.
- [17] Sumedh M. Joshi, Greg N. Thomsen, and Peter J. Diamessis. Deflation-accelerated preconditioning of the Poisson-Neumann Schur problem on long domains with a high-order discontinuous element-based collocation method. *Journal of Computational Physics*, page 24, 2016.
- [18] Marcello Manna, Andrea Vacca, and Michel O. Deville. Preconditioned spectral multi-domain discretization of the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 201:204–223, 2004.
- [19] L N Olson, J S Hesthaven, and L C Wilcox. Developments in Overlapping Schwarz Preconditioning of High-Order Nodal Discontinuous Galerkin Discretizations. In *Domain Decomposition Methods in Science and Engineering XVI*, number November 2004, pages 1–8. 2004.
- [20] Richard Pasquetti, Francesca Rapetti, Luca Pavarino, and Elena Zampieri. Neumann-Neumann-Schur complement methods for Fekete spectral elements. *Journal of Engineering Mathematics*, 56:323–335, 2006.
- [21] Luca F. Pavarino and Timothy Warburton. Overlapping Schwarz Methods for Unstructured Spectral Elements. *Journal of Computational Physics*, 160:298–317, 2000.
- [22] C. Pozrikidis. A Note on the Regularization of the Discrete Poisson-Neumann Problem. *Journal of Computational Physics*, 172:917–923, 2001.
- [23] I Yamazaki and X S Li. On techniques to improve robustness and scalability of the Schur complement method. In *Proceedings of the 9th International VECPAR Conference*, number 2, 2010.
- [24] I. Yamazaki, X. S. Li, and E. G. Ng. Preconditioning Schur complement systems of highly-indefinite linear systems for a parallel hybrid solver. *Numerical Mathematics*, 3:352–366, 2010.